February 12, 2003

Eric Rasmusen, Erasmuse@indiana.edu

10 Mechanism Design

10.1 The Revelation Principle and Moral Hazard with Hidden Knowledge Moral Hazard with Hidden Knowledge

Information is complete in moral hazard games, but in moral hazard with hidden knowledge the agent, but not the principal, observes a move of Nature after the game begins. Information is symmetric at the time of contracting but becomes asymmetric later. From the principal's point of view, agents are identical at the beginning of the game but develop private types midway through, depending on what they have seen. His chief concern is to give them incentives to disclose their types later, which gives games with hidden knowledge a flavor close to that of adverse selection. (In fact, an alternative name for this might be **post-contractual adverse selection.**) The agent might exert effort, but effort's contractibility is less important when the principal does not know which effort is appropriate because he is ignorant of the state of the world chosen by Nature. The main difference technically is that if information is symmetric at the start and only becomes asymmetric after a contract is signed, the participation constraint is based on the agent's expected payoffs across the different types of agent he might become. Thus, there is just one participation constraint even if there are eventually n possible types of agents in the model, rather than the n participation constraints that would be required in a standard adverse selection model.

There is more hope for obtaining efficient outcomes in moral hazard with hidden knowledge than in adverse selection or simple moral hazard. The advantage over adverse selection is that information is symmetric at the time of contracting, so neither player can use private information to extract surplus from the other by choosing inefficient contract terms. The advantage over simple moral hazard is that the post-contractual asymmetry is with respect to knowledge only, which is neutral in itself, rather than over whether the agent exerted high effort, which causes direct disutility to him.

Production Game VII: Hidden Knowledge

Players
The principal and the agent.

The Order of Play
1 The principal offers the agent a wage contract of the form $w(q, m)$, where q is output and m is a message to be sent by the agent.
2 The agent accepts or rejects the principal's offer.
3 Nature chooses the state of the world θ, according to probability distribution $F(\theta)$. The agent observes θ, but the principal does not.
4 If the agent accepts, he exerts effort e and sends a message m, both observed by the principal.
5 Output is $q(e, \theta)$.

Payoffs

If the agent rejects the contract, $\pi_{agent} = \bar{U}$ and $\pi_{principal} = 0$.
If the agent accepts the contract, $\pi_{agent} = U(e, w, \theta)$ and $\pi_{principal} = V(q - w)$.

The principal would like to know θ so he can tell which effort level is appropriate. In an ideal world he would employ an honest agent who always chose $m = \theta$, but in noncooperative games, talk is cheap. Since the agent's words are worthless, the principal must try to design a contract that either provides incentive for truth-telling or takes lying into account – he **implements** a **mechanism** to extract the agent's information.

The Revelation Principle

A principal might choose to offer a contract that induces his agent to lie in equilibrium, since he can take lying into account when he designs the contract, but this complicates the analysis. Each state of the world has a single truth, but a continuum of lies. Generically speaking, almost everything is false. The following principle helps us simplify contract design.

The Revelation Principle. *For every contract $w(q, m)$ that leads to lying (that is, to $m \neq \theta$), there is a contract $w^*(q, m)$ with the same outcome for every θ but no incentive for the agent to lie.*

Many possible contracts make false messages profitable for the agent because when the state of the world is a he receives a reward of x_1 for the true report of a and $x_2 > x_1$ for the false report of b. A contract which gives the agent the same reward of x_2 regardless of whether he reports a or b would lead to exactly the same payoffs for each player while giving the agent no incentive to lie. The revelation principle notes that a truth- telling contract like this can always be found by imitating the relation between states of the world and payoffs in the equilibrium of a contract with lying. The idea can also be applied to games in which two players must make reports to each other.

Applied to concrete examples, the revelation principle is obvious. Suppose we are concerned with the effect on the moral climate of cheating on income taxes, but anyone who makes $70,000 a year can claim he makes $50,000 and we do not have the resources to catch him. The revelation principle says that we can rewrite the tax code to set the tax to be the same for taxpayers earning \$70,000 and for those earning \$50,000, and the same amount of taxes will be collected without anyone having incentive to lie. Applied to moral education, the principle says that the mother who agrees never to punish her daughter if she tells her all her escapades will never hear any untruths. Clearly, the principle's usefulness is not so much to improve outcomes as to simplify contracts. The principal (and the modeller) need only look at contracts which induce truth-telling, so the relevant strategy space is shrunk and we can add a third constraint to the incentive compatibility and participation constraints to help calculate the equilibrium:

(3) **Truth-telling.** The equilibrium contract makes the agent willing to choose $m = \theta$.

The revelation principle says that a truth-telling equilibrium exists, but not that it is unique. It may well happen that the equilibrium is a weak Nash equilibrium in which the optimal contract gives the agent no incentive to lie but also no incentive to tell the truth. This is similar to the open-set problem discussed in Section 4.3; the optimal contract may satisfy the agent's participation constraint but makes him indifferent between accepting and rejecting the contract. If agents derive the slightest utility from telling the truth, of course, then truthtelling becomes a strong equilibrium, but if their utility from telling the truth is really significant, it should be made an explicit part of the model. If the utility of truth- telling is strong enough, in fact, agency problems and the costs associated with them disappear. This is one reason why morality is useful to business.

NOTES FOR CLASS

Handouts: old assignments.

In getting reference letters, editors are useful.

Also, this is a part of research, service to the profession. I do comments during seminars. It is satisfying and useful.

Getting comments is very hard. It is easier for students. People are not as grateful as they should be. It is a major use of co-authors– just to read the stuff.

These next few weeks: Some useful topics from my book. IOish topics. This week, mechanisms. Important for regulatoin, especially, but also price discrimination.

10.3: Myerson Mechanism Design Example

A seller has 100 units of a good.

If it is high quality, he values it at 40 dollars per unit; if it is low quality, at 20 dollars.

The buyer, who cannot observe quality before purchase, values high quality at 50 dollars per unit and low quality at 30 dollars.

The only way to get the seller to truthfully reveal the quality of the good, is for the buyer to say that if the seller admits the quality is bad, he will buy more units than if the seller claims it is good.

Who offers the contract?

When it is offered?

Myerson Trading Game I

Players

A buyer and a seller.

The Order of Play

1 The seller offers the buyer a contract $(Q_H, P_H, T_H, Q_L, P_L, T_L)$ under which the seller will later declare his quality to be high or low, and the buyer will first pays the lump sum T to the seller (perhaps with $T < 0$) and then buy Q units of the 100 the seller has available, at price P.
2 The buyer accepts or rejects the contract.
3 Nature chooses whether the seller's good is High quality (probability 0.2) or low quality (probability 0.8), unobserved by the buyer.
4. If the contract was accepted by both sides, the seller declares his type to be L or H and sells at the appropriate quantity and price as stated in the contract.

Payoffs

If the seller rejects the contract, $\pi_{buyer} = 0$ $\pi_{seller\ H} = 40 * 100$, and $\pi_{seller\ L} = 20 * 100$.
If the seller accepts the contract and declares a type that has price P and quantity Q, then

$$\pi_{buyer|seller\ H} = -T + (50 - P)Q \quad and \quad \pi_{buyer|seller\ L} = (30 - P)Q \quad (1)$$

and

$$\pi_{seller\ H} = 40(100 - Q) + PQ \quad and \quad \pi_{seller\ L} = 20(100 - Q) + PQ. \quad (2)$$

CONSTRAINTS

Participation Constraint:

$$
\begin{aligned}
0.8\pi_b(Q_L, P_L) + 0.2\pi_b(Q_H, P_H) &\geq 0 \\
0.8(30 - P_L)Q_L + 0.2(30 - P_H)Q_H &\geq 0
\end{aligned}
\tag{3}
$$

Incentive compatibility constraints, low quality:

$$
\begin{aligned}
\pi_L(Q_L, P_L) &\geq \pi_H(Q_H, P_H) \\
20(100 - Q_L) + P_L Q_L &\geq 20(100 - Q_H) + P_H Q_H,
\end{aligned}
\tag{4}
$$

High quality:

$$
\begin{aligned}
\pi_H(Q_H, P_H) &\geq \pi_H(Q_L, P_L) \\
30(100 - Q_H) + P_H Q_H &\geq 30(100 - Q_L) + P_L Q_L.
\end{aligned}
\tag{5}
$$

To make the contract incentive compatible, the seller needs to set P_H greater than P_L, but if he does that it will be necessary to set Q_H less than Q_L. If he does that, then the low-quality seller will not be irresistably tempted to pretend his quality is high: he would be able to sell at a higher price, but not as great a quantity.

Since Q_H is being set below 100 only to make pretending to be high-quality unattractive, there is no reason to set Q_L below 100, so $Q_L = 100$. The buyer will accept the contract if $P_L \leq 30$, so the seller should set $P_L = 30$.

Myerson Trading Game I, continued

The low-quality seller's incentive compatibility constraint, inequality (4), will be binding, and thus becomes

$$\begin{aligned} \pi_L(Q_L, P_L) &\geq \pi_H(Q_H, P_H) \\ 20(100-100) + 30 * 100 &= 20(100 - Q_H) + P_H Q_H. \end{aligned} \tag{6}$$

Solving for Q_H gives us $Q_H = \frac{1000}{P_H - 20}$, which when substituted into the seller's payoff function yields

$$\begin{aligned} \pi_s &= 0.8\pi_L(Q_L, P_L) + 0.2\pi_H(Q_H, P_H) \\ &= 0.8[(20)(100 - Q_L) + P_L Q_L] + 0.2[(40)(100 - Q_H) + P_H Q_H] \\ &= 0.8[(20)(100-100) + 30 * 100] + 0.2[(40)(100 - \tfrac{1000}{P_H-20}) + \\ & P_H(\tfrac{1000}{P_H-20})] \end{aligned} \tag{7}$$

Maximizing with respect to P_H subject to the constraint that $P_H \leq 50$ (or else the buyer will turn down the contract) yields the corner solution of $P_H = 50$, which allows for $Q_H = 33\frac{1}{3}$.

The participation constraint for the buyer is already binding, so we do not need the transfers T_L and T_H to take away any remaining surplus, as we might in other situations.[1] Thus, the equilibrium contract is

$$(Q_L = 100, P_L = 30, T_L = 0, Q_H = 33\frac{1}{3}, P_H = 50, T_H = 0). \tag{8}$$

[1] The transfers could be used to adjust the prices, too. We could have $Q_L = 20$ and $T_L = 1000$ in equation (7) without changing anything important.

Myerson Trading Game II

The Order of Play: The same as in Myerson Trading Game I except that the buyer makes the contract offer in move (1) and the seller accepts or rejects in move (2).

Payoffs: The same as in Myerson Trading Game I.

The participation constraint in the buyer's mechanism design problem is

$$0.8\pi_L(Q_L, P_L) + 0.2\pi_H(Q_H, P_H) \geq 0. \tag{9}$$

The incentive compatibility constraints are just as they were before.

As before, the mechanism will set $Q_L = 100$, but it will have to make $Q_H < 100$ to deter the low-quality seller from pretending he is high- quality. Also, $P_H \geq 40$, or the high-quality seller will pretend to be low-quality.

Suppose $P_H = 40$. The low-quality seller's incentive compatibility constraint will be binding, and thus becomes

$$\begin{aligned} \pi_L(Q_L, P_L) &\geq \pi_H(Q_H, P_H) \\ 20(100 - 100) + P_L * 100 &= 20(100 - Q_H) + 40Q_H. \end{aligned} \tag{10}$$

Myerson Trading Game II, continued

Solving for Q_H gives us $Q_H = 5P_L - 100$, which when substituted into the buyer's payoff function yields

$$\begin{aligned}
\pi_b &= 0.8\pi_{b|L}(Q_L, P_L) + 0.2\pi_{b|H}(Q_H, P_H) \\
&= 0.8[(30 - P_L)Q_L] + 0.2[(50 - P_H)Q_H] \\
&= 0.8[(30 - P_L)100] + 0.2[(50 - 40)(5P_L - 100)] \\
&= 2400 - 80P_L + 10P_L - 200 = 2200 - 70P_L
\end{aligned} \tag{11}$$

Maximizing with respect to P_L subject to the constraint that $P_L \geq 20$ (or else we would come out with $Q_H < 0$ to satisfy incentive compatibility constraint (6)) yields the corner solution of $P_L = 20$, which requires that $Q_H = 0$.

Would setting $P_H > 40$ help? No, because that just makes it harder to satisfy the low-quality seller's incentive compatibility constraint. We would continue to have $Q_H = 0$, and, of course, P_H does not matter if nothing is sold. And as before, we do not need to make use of transfers to make the participation constraint binding. Thus, the equilibrium contract has P_H take any possible value and

$$(Q_L = 100, P_L = 20, T_L = 0, Q_H = 0, T_H = 0). \tag{12}$$

Myerson Trading Game III

The Order of Play

0. Nature chooses whether the seller's good is high quality (probability 0.2) or low quality (probability 0.8), unobserved by the buyer.
1 The buyer offers the seller a contract $(Q_H, P_H, T_H, Q_L, P_L, T_L)$ under which the seller will later declare his quality to be high or low, and the buyer will first pays the lump sum T to the seller (perhaps with $T < 0$) and then buy Q units of the 100 the seller has available, at price P.
2 The seller accepts or rejects the contract.
3. If the contract was accepted by both sides, the seller declares his type to be L or H and sells at the appropriate quantity and price as stated in the contract.

Payoffs: The same as in Myerson Trading Games I and II.

Myerson Trading Game III, continued

The incentive compatibility constraints are unchanged from the previous two versions of the game, but now the participation constraints are different for the two types of seller.

$$\pi_L(Q_L, P_L) \geq 0 \tag{13}$$

and

$$\pi_H(Q_H, P_H) \geq 0. \tag{14}$$

Any mechanism which satisfies these two constraints would also satisfy the single participation constraint in MTG II, since it says that a weighted average of the payoffs of the two sellers must be positive. Thus, any mechanism which maximized the buyer's payoff in MTG II would also maximize his payoff in MTG III, if it satisfied the tougher bifurcated participation constraints. The mechanism we found for the game does satisfy the tougher constraints, so it is the optimal mechanism here too.

This is not a general feature of mechanisms. More generally the optimal mechanism will not have as high a payoff when one player starts the game with superior information, because of the extra constraints on the mechanism.

Myerson Trading Game IV

The Order of Play: The same as in Myerson Trading Game III except that in (1) the seller makes the offer and in (2) the buyer accepts or rejects.

Payoffs: The same as in Myerson Trading Games I, II, and III.

The incentive compatibility constraints are the same as in the previous games, and the participation constraint is inequality (3), just as in Myerson Trading Game I. The big difference now is that unlike in the first three versions, MTG IV has an informed player making the contract offer. As a result, the form of the offer can convey information, and we have to consider out-of- equilibrium beliefs, as in the dynamic games of incomplete information in Chapter 6 (and we will see more of this in the signalling models of Chapter 11). Surprisingly, however, the importance of out-of-equilibrium beliefs does not lead to multiple equilibria. Instead, the equilibrium contract is

M1: $(Q_L = 100, P_L = 30, T_L = 0, Q_H = 33\frac{1}{3}, P_H = 50, T_H = 0)$,

This is part of equilibrium under the out-of-equilibrium belief that if the seller offers any other contract, the buyer believes the quality is low.

This is the same equilibrium mechanism as in MTG I.

MTG IV, continued

Consider two other mechanisms, M2 and M3 which satisfy the two incentive compatibility constraints and the participation constraint, but which are not equilibrium choices:

M2: $(Q_L = 100, P_L = 28, T_L = 0, Q_H = 0, P_H = 40, T_H = 800)$.

M3: $(Q_L = 100, P_L = 31\frac{3}{7}, T_L = 0, Q_H = 57\frac{1}{7}, P_H = 40, T_H = 0)$.

Mechanism M2 is interesting because the buyer expects a positive payoff of (30- 28) (100) = 200 if the seller is low-quality and a negative payoff of 800 if the seller is high-quality, for an overall expected payoff of zero. The contract is incentive compatible because a low-quality seller could not increase his payoff of 28*100 by pretending to be high-quality (he would get 20*100 + 800 instead), and a high-quality seller would reduce his payoff of (40*100 + 800) if he pretended to have low quality. Here, for the first time, we see a positive value for the transfer T_H.

Under mechanism M3, the buyer expects a negative payoff of (30- 31 $\frac{1}{7}$) (100) = -11 $\frac{3}{7}$ if the seller is low- quality and a positive payoff of $(57\frac{1}{7})(50 - 40)$ =11 $\frac{3}{7}$ if the seller is high-quality, for an overall expected payoff of zero. The contract is incentive compatible because a low- quality seller could not increase his payoff of 3,142 $\frac{6}{7}$ (31 $\frac{3}{7}$) (100) by pretending to be high-quality (he would get (57 $\frac{1}{7}$) (40) + (42 $\frac{6}{7}$) (20) instead, which comes to the same figure), and a high-quality seller would reduce his payoff if he pretended to have low quality and sold something he valued at 40 at a price of 31 $\frac{1}{7}$.

What Contract Does the Seller Like?

M1: $(Q_L = 100, P_L = 30, T_L = 0, Q_H = 33\frac{1}{3}, P_H = 50, T_H = 0)$,

M2: $(Q_L = 100, P_L = 28, T_L = 0, Q_H = 0, P_H = 40, T_H = 800)$.

M3: $(Q_L = 100, P_L = 31\frac{3}{7}, T_L = 0, Q_H = 57\frac{1}{7}, P_H = 40, T_H = 0)$.

Mechanism M1 maximizes the payoff of the average seller, as we found in MTG I, yielding the low-quality seller a payoff of 3,000 and the high-quality seller a payoff of 4,333 ($= (33\frac{1}{3})(50) + 66\frac{2}{3}(40)$), for an average payoff of 3,867. If the seller is high-quality, however, he would prefer mechanism M2, which has payoffs of 2800 and 4800 (=800+ 40(100)), for an average payoff of 3200. If the seller is low-quality, he would prefer mechanism M3, which has payoffs of $3,142\frac{6}{7}$ and 4000, for an average payoff of $3,314\frac{6}{7}$.

Suppose that the seller chose M2, regardless of his type. This could not be an equilibrium, because a low-quality seller would want to deviate. Suppose he deviated and offered a contract almost like M1, except that $P_L = 29.99$ instead of 30 and $P_H = 49.99$ instead of 50. This new contract would yield positive expected payoff to the buyer whether the buyer believes the seller is low-quality or high-quality, and so it would be accepted. It would yield higher payoff to the low- quality seller than M2, and so the deviation would have been profitable. If the seller chose M3 regardless of his type, a high- quality seller could profitably deviate in the same way.